CLAIMS

1.   Process for assigning tasks in a multiprocessor digital data processing system with a preemptive operating system, comprising a given number of processors capable of processing said tasks in parallel, characterized in that it comprises at least one preliminary phase during which said processors (20a-21a, 20b-22b, 20c) are divided into groups (Ga, Gb, Gc), each group comprising predetermined numbers of processors, in that each of said processor groups is associated with an elementary queue (5a, 5b, 5c) storing a predetermined number of tasks to be processed in a given order of priority, and in that each of the tasks is associated with one of the processors associated with this elementary queue (5a, 5b, 5c).

2.   Process according to claim 1, characterized in that said groups each comprise an identical number of processors (200-203, 210-213).

3.   Process according to claim 1 or 2, characterized in that it comprises an additional preliminary phase consisting of generating a series of tests and measurements for determining the number of processors in each group and the number of groups that make it possible to achieve the best performance of said system.

4.   Process according to any of claims 1 through 3, characterized in that, the architecture of said system being of the non-uniform memory access type known as "NUMA," and the system (1) being constituted by a predetermined number of modules (M0, M1) linked to one another, each comprising a given number of processors (200-203, 210-213) and storage means, each of said modules (M0, M1) constitutes one of said groups, each module being associated with one of said elementary queues.

5.   Process according to any of claims 1 through 4, characterized in that each of said processors is associated with a first data structure that identifies it, in that said first data structure comprises at least one first set of pointers (p200 through p203) associating it with one of said

4   elementary queues (5a, 5b), in that each of said elementary queues (5a, 5b) is associated with a

5   second data structure, in that said second data structure comprises at least one second set of

6   pointers (pp5a, pp5b) associating it with one of said processor groups (200-201, 202-203), in that

7   all of the tasks to be processed (T1 through T10) in said system (1) being stored in a table (4),

8   each of said second data structures of the elementary queues (5a, 5b) also comprises a third set of

9   pointers (pT1, pT5, pT10), each associating elementary queues (5a, 5b) with one of said tasks (T1

10   through T10) stored in the table (4) or with a series of concatenated tasks, and in that each of said

11   tasks (T1 through T10) of the table (4) is associated with a third data structure that comprises a

12   fourth set of pointers (p5a1 through p5a4, p5b1 through p5b10) associating it with one of said

13   elementary queues (5a, 5b).

1   6.   Process according to any of claims 1 through 5, characterized in that it comprises

2   at least one additional phase consisting of distributing said tasks among said elementary queues

3   (5a, 5b) by searching, when a new task to be processed (Tz) is created, for the queue with the

4   lightest load (5y) among all of said elementary queues (5a, 5x, 5y, 5p) of said system (1) and of

5   assigning said new task to this elementary queue so as to balance the global load of this system

6   (1) among said elementary queues (5a, 5x, 5y, 5p).

1   7.   Process according to claim 6, characterized in that said distribution is performed

2   by determining a so-called composite load parameter associated with each of said elementary

3   queues (5a, 5x, 5y, 5p) and in that, each processor (2a, 2x, 2y, 2p) being associated with storage

4   means (Mema, Memx, Memy, Memp), said composite load parameter is calculated as being the

5   sum of the load of a processor or a processor group associated with said elementary queue and

6   the load of the storage means associated with this processor or processor group.

1   8.   Process according to claim 6, characterized in that it comprises a preliminary step

2   consisting of checking whether said task (Tz) is linked to one of said elementary queues (5a, 5x,

3   5y, 5p), and when said test is positive, of assigning said new task to this elementary queue.

1   9.   Process according to any of claims 1 through 5, characterized in that it comprises

2   at least one additional phase consisting, when one of said elementary queues (5q) associated with

3   one of said processor groups (2q) is empty of executable tasks, of searching for a so-called

4   remote elementary queue (5y) that is not empty, and of selecting in this elementary queue (5y) a

5   task executable by one of said processors (2q) of said processor group associated with the empty

6   elementary queue (5q) and of transmitting it to this processor (2q) to be processed by it, so as to

7   globally balance the processing of said tasks in said system (1).

1   10.   Process according to claim 9, characterized in that said non-empty elementary

2   queue (5y) must have a predetermined minimal occupation threshold.

1   11.   Process according to claim 10, characterized in that furthermore, the tasks being

2   stored in decreasing order of priority, a predetermined number of tasks is skipped before

3   scanning the other tasks of said non-empty elementary queue (5y) in order to search for an

4   executable task and have it processed by one of said processors (2q) of said processor group

5   associated with the empty elementary queue (5q).

1   12.   Process according to claim 11, characterized in that said number of skipped tasks

2   and the maximum number of scanned tasks among all of those stored in said non-empty

3   elementary queue (5q) are variable over time and are determined by a self-adapting process from

4   the number of tasks that are or are not found during said scans and from the position of these

5   tasks, sequenced in order of priority, in said non-empty elementary queue (5q).

1   13.   Process according to any of claims 9 through 12, characterized in that said

2   selected task is that associated with a minimal value of a so-called cost parameter, which

3   measures the global performance degradation of said system (1) due to the processing of said

4   selected task in said non-empty remote elementary queue (5q) by one of said processors of said

5   processor group associated with the empty elementary queue (2q).

1   14.   Process according to any of claims 1 through 5, characterized in that it comprises

2   at least one additional phase comprising at least one step for periodically measuring for a

3   balanced distribution of said tasks in said elementary queues (5a, 5x, 5y, 5p) and, when an

4   unbalanced state of said system (1) is determined, a step for selectively moving tasks from at

5   least one elementary queue with a heavier load (5x) to an elementary queue with a lighter load

6   (5y).

1   15.   Process according to claim 14, characterized in that when said imbalance is below

2   a certain threshold, no moving of tasks is performed.

1   16.   Process according to claim 14 or 15, characterized in that, all or some of said

2   tasks belonging to multitask processes, each multitask process requiring a given memory size

3   and workload, it comprises a step for measuring said workloads and said memory sizes, in that it

4   comprises the selection of the process requiring the greatest workload and the smallest memory

5   size, and in that all the tasks of said selected process are moved to the elementary queue with the

6   lightest load (5y).

1   17.   Process according to claim 16, characterized in that it comprises a preliminary

2   step consisting of checking whether all of the tasks of said multitask process that must be moved

3   belong to the elementary queue set with the heaviest load (5x) and whether any task is linked to

4   any of said groups.

1   18.   Process according to any of claims 1 through 17, characterized in that said

2   operating system is of the "UNIX" (registered trademark) type.

1   19.   Architecture for a multiprocessor digital data processing system comprising a

2   given number of processors for implementing the process for assigning tasks to be processed to

3   said processors according to any of claims 1 through 18, characterized in that said processors

4   (20a-21a, 20b-22b, 20c) are divided into groups (Ga, Gb, Gc), and in that an elementary queue

5   (5a, 5b, 5c) associated with each of the groups (Ga, Gb, Gc) is provided, each of said elementary

6   queues (5a, 5b, 5c) storing a predetermined number of tasks to be processed in a given order of

7   priority, so that each of the tasks of each of said elementary queues (5a, 5b, 5c) is associated with

8   one of the processors of this elementary queue (20a-21a, 20b-22b, 20c).

1   20.   Architecture according to claim 19, characterized in that it also comprises means

2   (6) for determining the load of said elementary queues (5a, 5x, 5y, 5p) and for assigning a new

3   task created in said system to the elementary queue with the lightest load (5y).

1   21.   Architecture according to claim 19, characterized in that it also comprises, when

2   one (5q) of said elementary queues (5a, 5x, 5y, 5p) associated with one of said processors (2q) is

3   empty, means (7) for finding a non-empty, so-called remote elementary queue (5y), and for

4   finding an executable task in this elementary queue (5y), and assigning it to said processor (2q)

5   for processing.

1   22.   Architecture according to claim 19, characterized in that it also comprises means

2   (8) for detecting an imbalance between elementary queues (5a, 5x, 5y, 5p), and when such an

3   imbalance is detected, for determining the elementary queue with the heaviest load (5x) and the

4   elementary queue with the lightest load (5y), and for moving tasks from the elementary queue

5   with the heaviest load (5x) to the elementary queue with the lightest load (5y).

1   23.   Architecture according to any of claims 19 through 22, characterized in that, being

2   of the nonuniform memory access type known as "NUMA," composed of modules (M0, M1)

3   linked to one another, each comprising a given number of processors (200-203, 210-213) and

4   storage means, each of said modules (M0, M1) constitutes one of said groups, each module (M0,

5   M1) being associated with one of said elementary queues.